arXiv:1401.7893v1 [math.ST] 30 Jan 2014

Inference with penalized likelihood

Daniel Commenges [1,2] Jérémie Bureau [1,2] Hein Putter [3]

January 31, 2014

[1] INSERM, ISPED, Centre INSERM U-897-Epidemiologie-Biostatistique, Bordeaux, F-33000

[2] Univ. Bordeaux, ISPED, Centre INSERM U-897-Epidemiologie-Biostatistique, Bordeaux, F-33000, France

[3] Univ. Leiden, The Netherlands

ABSTRACT.

This work studies the statistical properties of the maximum penalized likelihood approach in a semi-parametric framework. We recall the penalized likelihood approach for estimating a function and review some asymptotic results. We investigate the properties of two estimators of the variance of maximum penalized likelihood estimators: sandwich estimator and a Bayesian estimator. The coverage rates of confidence intervals based on these estimators are studied through a simulation study of survival data. In a first simulation the coverage rates for the survival function and the hazard function are evaluated. In a second simulation data are generated from a proportional hazard model with covariates. The estimators of the variances of the regression coefficients are studied. As for the survival and hazard functions, both sandwich and Bayesian estimators exhibit relatively good properties, but the Bayesian estimator seems to be more accurate. As for the regression coefficients, we focussed on the Bayesian estimator and found that it yielded good coverage rates.

Key Words : Penalized likelihood; crossvalidation; sandwich estimator; splines; asymptotics; survival data

1 Introduction

Penalized likelihood has been widely used for non or semi-parametric estimation of a function. It is also used in parametric models for variable selection. Here we are interested in its use for estimating a hazard function as proposed by O'Sullivan (1986) and Joly et al. (1998). More generally, this approach have been used for estimating transition intensities in multi-state models Joly et al. (2002). It has the advantage that while making no parametric assumption on the hazard or intensity functions, it yields smooth estimates of these functions. Moreover, complex cases like interval-censored observations can be treated easily, at least conceptually. The main difficulty in the approach is the choice of the smoothing coefficient that can been done by cross-validation. There is also a theoretical difficulty in studying the properties of the maximum penalized likelihood estimators. There are mainly two candidates for estimating the variance of the maximum penalized likelihood estimators, inspired by M-estimators and Bayesian theory respectively. However precise theoretical results are lacking for justifying these estimators of the variance. The aim of this paper is to recall some facts about penalized likelihood and to study the two approaches for estimating the variances by simulation.

Section 2 recalls the penalized likelihood approach for obtaining smooth non- or semi-parametric estimates of functions, together with the choice of the smoothing coefficient by approximate cross-validation. In section 3 we recall some asymptotic results and we investigate possible estimators of the variance of the penalized likelihood estimators, especially using a Bayesian approach and results from M-estimators theory. The case of survival data is developed in section 4. A simulation study is presented in section 5. Section 6 concludes.

2 The penalized likelihood and crossvalidation

2.1 Penalized likelihood

In the case where at least one of the parameters is a function, penalized likelihood can be applied for obtaining a smooth estimator (O'Sullivan, 1986). The penalized loglikelihood is

$$\mathrm{pL}(\theta;\kappa) = L(\theta) - \kappa J(\theta),$$

where $L(\theta)$ is the loglikelihood and $\kappa J(\theta)$ is a penalty term. For instance if the sample $\bar{\mathcal{O}}_n$ is constituted of n i.i.d. variables Y_i, the loglikelihood for g^θ is $\sum_{i=1}^n \log g^\theta(Y_i)$. We consider the semi-parametric case where $\theta = (\alpha, \beta)$ where α is a function and β is a

real parameter. In this case, the penalty term often depends only on the function α. The most often used form of the penalty is $J(\alpha) = \int [\alpha''(t)]^2 \, dt$. The smoothing parameter κ weighs the penalty term. The maximum penalized likelihood estimator (MPLE) is $\hat{\theta}_\kappa$ which maximizes $\mathrm{pL}(\theta; \kappa)$.

The MPLE of the function is most often approximated on a spline basis so that the problem reduces to a parametric one with a large number of parameters. In several cases, a spline is indeed the exact solution (Wahba, 1983), so that a finite number of real parameters have to be estimated. Even in that case, the model is still nonparametric because the location of the knots is driven by the observations. In general, the spline is used for approximating the solution. In theory, the number of knots can be arbitrarily large: the larger, the better the approximation of $\hat{\theta}_\kappa$. In the penalized likelihood approach, the splines are used only for approximating the solution of the maximisation of $\mathrm{pL}(\theta; \kappa)$: the degree of smoothness is tuned by the weight κ given to the penalty.

2.2 The approximate crossvalidation criterion

The choice of the smoothing coefficient κ is crucial. It is desirable that this choice be data-driven and the most common approach is cross-validation. We will focus on likelihood cross-validation which estimates the crossentropy of the estimator (Commenges et al., 2012). The leave-one-out cross-validation criterion may be computationally demanding since it is necessary to run the maximization algorithm n times for finding the $\hat{\theta}_{-i}, i = 1, \ldots, n$. For this reason an approximate formula is very useful. Commenges et al. (2007) gave the following formula:

$$\mathrm{LCV}_a(\kappa) = -n^{-1} L(\hat{\theta}_\kappa) + \mathrm{Trace}(H_{\mathrm{pL}}^{-1} K), \tag{1}$$

where $K = n^{-1} \sum_{i=1}^{n} \hat{v}_i \hat{d}_i^T$ with $\hat{v}_i = \frac{\partial L_i(\theta)}{\partial \theta}|_{\hat{\theta}_\kappa}$ the individual score taken at the value of the penalized likelihood estimator, and $\hat{d}_i = \frac{1}{n-1}(\hat{v}_i + \kappa \frac{\partial J(\theta)}{\partial \theta}|_{\hat{\theta}_\kappa})$, and where $H_{\mathrm{pL}} = -\frac{\partial^2 L(\theta)}{\partial \theta^2}|_{\hat{\theta}_\kappa} + \kappa \frac{\partial^2 J(\theta)}{\partial \theta^2}|_{\hat{\theta}_\kappa}$. The correction term is often interpreted as the equivalent number of parameters (Commenges et al., 2007).

3 Asymptotic distribution of the penalized likelihood estimators

3.1 Consistency

Cox and O'Sullivan (1990, 1996) present consistency results for MPLE in a non parametric framework. The problem is formalized using a normalized penalized loglikelihood. We define $p\bar{L}_n = \bar{L}_n(\theta) - \lambda_n J(\theta)$ where $\bar{L}_n = n^{-1}L(\theta)$ and $\lambda_n = n^{-1}\kappa_n$. For consistency of $\hat{\theta}_n$, the sequence λ_n must tend toward zero in a certain range of speed. We note $\lambda_n = O(n^{-\alpha})$, and consider $0 \leq \alpha \leq 1$. If λ_n tends to 0 when $n \to \infty$ the bias of $\hat{\theta}_n$ goes to 0. The variance of $\hat{\theta}_n$ tends to zero whatever α, $0 \leq \alpha \leq 1$. Consistency obtains whatever $\alpha > 0$ (Yu and Ruppert, 2002).

In the parametric case we can deduce the consistency of the MPLE from that of the MLE in the case of $\lambda_n = o(n)$.

Proposition 3.1. *Let Θ be a compact parameter space. We assume that the Maximum Likelihood Estimator (MLE) $\hat{\theta}^{ML}$ satisfies the conditions of the non penalized maximum likelihood estimator of the theorem 5.7 from (Van der Vaart, 2000) : there exists a fixed function which not depend on n $L(\theta; x) = \mathbb{E}[L_n(\theta; x)]$, and a value $\theta_0 \in \Theta$ such that, $\forall \varepsilon > 0$,*

$$\sup_{\theta \in \Theta} |\bar{L}_n(\theta; x) - L(\theta)| \xrightarrow{p} 0 \quad \textit{(uniform convergence)} \tag{2}$$

$$\sup_{\theta : \|\theta - \theta_0\| \leq \varepsilon} L(\theta) < L(\theta_0) \tag{3}$$

Then, if $\lambda_n = o(n)$, $\hat{\theta}^{pL}$ converges in probability to θ_0.

Proof. Let $\theta \in \Theta$, the parameter space defined as a compact.

$$\sup_{\theta \in \Theta} |p\bar{L}_n(\theta) - L(\theta)| \leq \sup_{\theta \in \Theta} |\lambda_n J(\theta)| + \sup_{\theta \in \Theta} |\bar{L}_n(\theta) - L(\theta)| \tag{4}$$

if $\lambda_n = o(n)$ then,

$$\sup_{\theta \in \Theta} |\lambda_n J(\theta)| \longrightarrow 0 \quad \text{as } n \to \infty. \tag{5}$$

And by the classical non penalized MLE properties,

$$\sup_{\theta \in \Theta} |\bar{L}_n(\theta; x) - L(\theta)| \xrightarrow{p} 0 \quad \text{as } n \to \infty \tag{6}$$

The function $L(\theta)$ satisfies (MLE properties) the root unicity condition : let $\theta_0 \in \Theta$ such that

$$\sup_{\theta : \|\theta - \theta_0\| \leq \varepsilon} L(\theta) < L(\theta_0) \tag{7}$$

According to the theorem 5.7 of (Van der Vaart, 2000), $\hat{\theta}^{pL} \xrightarrow{p} \theta_0$ as $n \to \infty$. □

The variance increases while the bias decreases with α. There is an optimal rate for the MSE which depends on the degree of differentiability of the unknown function. It is n^{-1} if the degree of differentiability of the unknown function is infinite. As we shall see in section 3.2 there is a link between penalized likelihood and the Bayesian approach. The penalized log-likelihood can be interpreted as the logarithm of the numerator of the posterior distribution and the term $\kappa J(\theta)$ is then up to an additive constant the log of the prior. With this interpretation, since the prior is fixed, κ does not depend on n, and thus $\lambda_n = O(n^{-1})$.

In spite of their theoretical interest, these results are not very useful in practice where the problem is to find the optimal value for a given n. The most often used method for doing this is crossvalidation. There are results in density estimation about the optimality of the crossvalidation choice: the estimator using the value of λ_n chosen by crossvalidation has asymptotically the same rate of convergence than using the optimal value of λ_n (Hall et al., 1992).

One may wonder how the sequence of λ_n chosen by crossvalidation behaves. We performed a simulation of survival data following a Weibull distribution and with samples of different sizes going from 100 to 2000 observations. The smoothing parameters κ_n were estimated with the approximate cross validation criterion LCV_a. Figure 1 shows the behavior of the sequence κ_n, $n^{-1/2}\kappa_n$ and $\lambda_n = n^{-1}\kappa_n$ as a function of n. Apparently the sequence λ_n goes to 0 as the number of observations grows and the rate seems to slightly slower than n^{-1}, the optimal rate for infinitely differentiable functions.

3.2 Penalized likelihood as a posterior distribution

Wahba (1983) has made the link between penalized likelihood and the Bayesian approach. We must consider probability measures on a functional space. In this approach θ is considered as a random variable and the penalized loglikelihood appears as the numerator of the posterior density, where $\kappa J(\theta)$ is the log-prior density (up to an additive constant). Either κ is known, or it is an hyper-parameter. It can be estimated from the data in an empirical Bayes spirit. Anyway, as a parameter it should not depend on n. Since $\lambda_n = n^{-1}\kappa_n$, this choice is compatible with the optimal order of the sequence λ_n in n^{-1} obtained by Cox and O'Sullivan (1996) in the case of infinite differentiability.

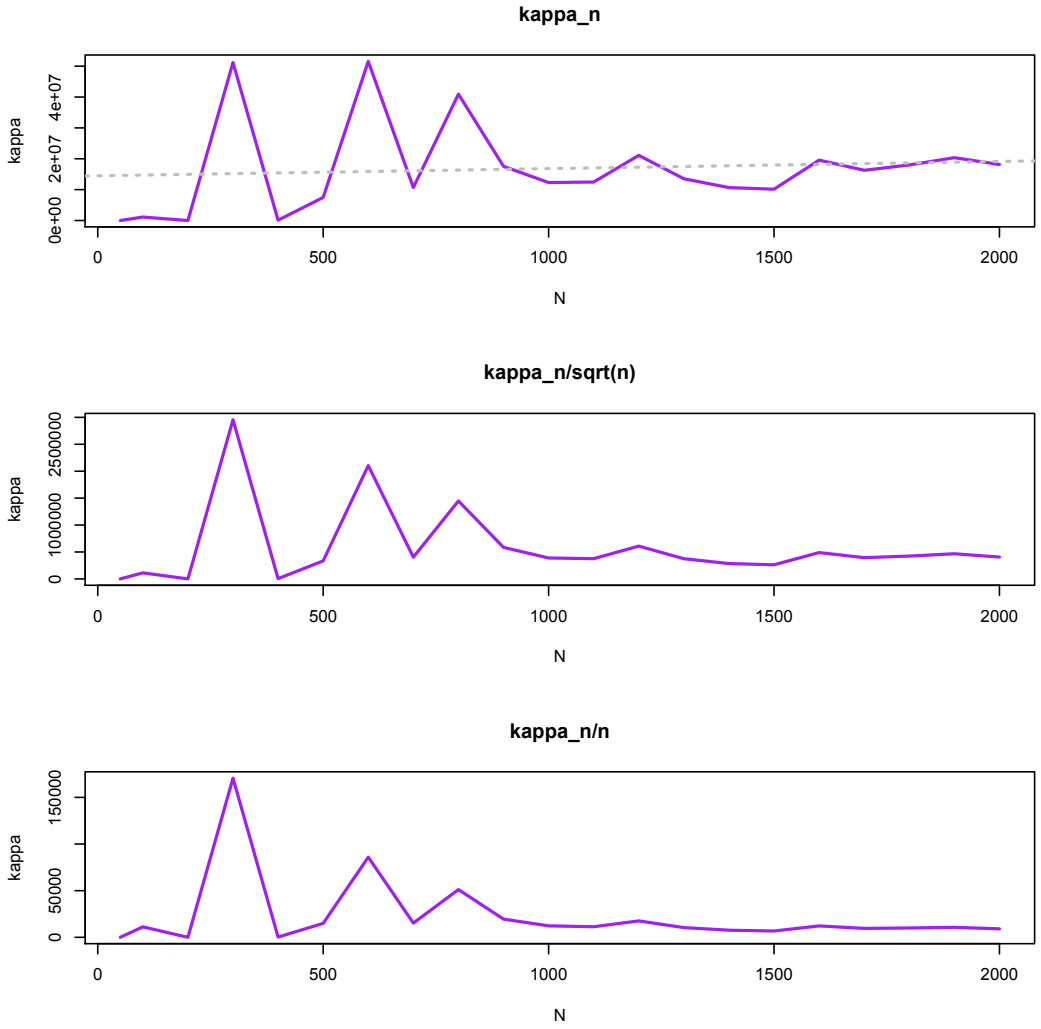

Figure 1: Sequences κ_n, $n^{-1/2}\kappa_n$ and λ_n. The sequence of smoothing parameter κ_n is computed using the LCV_a crossvalidation criterion.

O'Sullivan (1988) proposed an estimator based on a Bayesian view of the penalized likelihood. This leads to approximate the distribution of $\hat{\theta}$ by a Gaussian distribution with variance $V_n^{\text{Bayes}}(\hat{\theta}) = H_{pL_n}^{-1}(\hat{\theta})$, where $H_{pL_n}(\hat{\theta}) = -\frac{\partial^2}{\partial\theta^2} pL_n(\hat{\theta})$, the Hessian of minus the logarithm of the penalized likelihood. This estimator will be called the Bayesian estimator.

3.3 Penalized likelihood estimators as M-estimators

Since Huber's papers in the 1960's, M-estimation methods became important for asymptotic analysis and approximate inference. Huber (1964) introduced M-estimators and their asymptotic properties and they were an important part for the development of robust statistics. M-estimators were used to sudy the properties of maximum likelihood estimators in misspecified models. Zeger and Liang (1986) used M-estimator theory to justify the use of generalized estimating equations in the case of correlated observations.

Assuming $Y_i, \cdots, Y_n$ are independent and identically distributed, we can consider that the MPLE is a M-estimator (Van der Vaart, 2000), that is an estimator which satisfies

$$n^{-1}\sum_{i=1}^{n}\phi(\theta, Y_i) = 0 \quad \text{, with} \quad \phi(\theta, Y_i) = -\frac{\partial L_i(\theta)}{\partial\theta} + \kappa\frac{\partial J(\theta)}{\partial\theta}. \tag{8}$$

The asymptotic distribution of $\hat{\theta}$ is the multivariate normal distribution $\mathcal{N}(\theta_0, V(\theta_0))$, where θ_0 is the parameter value defined by $\mathrm{E}[\phi(\theta_0, Y_1)] = 0$ and where V can be estimated by the (penalized) sandwich estimator :

$$\hat{V}(\theta_0) = H_{\mathrm{pL}}^{-1}(\theta_0)\left[\sum U_i U_i^T\right]H_{\mathrm{pL}}^{-1}(\theta_0), \tag{9}$$

where the individual penalized score function are given by $U_i = (n-1)\hat{d}_i$, $\hat{d}_i = \frac{1}{n-1}(\hat{v}_i + \kappa\frac{\partial J(\theta)}{\partial\theta}|_{\theta_0})$ and the non penalized score function $\hat{v}_i = \frac{\partial L_i(\theta)}{\partial\theta}|_{\theta_0}$.

We also define the non-penalized sandwich estimator obtained by putting $\kappa = 0$ in the formula. In that case, the non-penalized scores and Hessian are used and the formula is similar to the robust variance estimator of the variance of the MLE, except that the parameter value is taken at the maximum of the penalized likelihood (rather than the maximum of the likelihood).

M-estimation theory needs that the function ϕ does not depend on i and n. Chen and Zhang (2009) used a M-estimator for functional estimation with a function depending on n through a smoothing coefficient represented by a window width in a kernel method. Cheng and Huang (2010) also derive results in which the $\phi(\theta, Y_i)$ may depend on n and proves the consistency of a distribution obtained by bootstrap. Indeed bootstrap is another way of doing inference; however it is not feasible in complex problems, especially those in which the loglikelihood involves numerical integrals (as is the case for interval-censored observations). Actually, Stefanski and Boos (2002) gave an extension of the M-estimators theory in which θ satisfies $\sum_{i=1}^{n}\phi(\theta, Y_i) = s_n$, where $s_n/\sqrt{n} \xrightarrow{p} 0$ as $n \to \infty$. In this particular case, the asymptotic property of $\hat{\theta}$ holds. This extension allows to cover a wider class of statistics whose ϕ function depends on n.

4 Penalized likelihood inference for survival analysis

We will focus on a proportional hazard model for survival data. We denote the independent survival times by $T_1, \cdots, T_n$. For $i = 1, \ldots, n$ we observe $\tilde{T}_i = \min(T_i, C_i)$, where C_i is a censoring variable. We use $i \in \mathcal{O}$ to denote the indexes of uncensored survival times. The proportional hazard model is given by the following expression for the hazard function:

$$h_i(t) = h_0(t) \exp(X_i \beta) \tag{10}$$

The aim of the study is to estimate the baseline hazard function $h_0(t)$ and the p parameters vector β of the covariates X_i.

The full likelihood function is then given by :

$$L(\beta, h_0) = -\sum_{i=1}^{n} H_0(\tilde{T}_i) \exp(X_i \beta) + \sum_{i \in \mathcal{O}} (\log h_0(\tilde{T}_i) + X_i \beta) \tag{11}$$

where $H_0(t) = \int_0^t h_0(u)du$ is the cumulative baseline hazard function. We want to estimate the couple (β, h_0) by using a penalized likelihood approach.

The penalized likelihood function is defined by

$$pL(\beta, h_0) = L(\beta, h_0) - \kappa J(h_0) \quad , \ \kappa \geq 0. \tag{12}$$

where $J(h_0)$ is the penalty function defined by $J(h_0) = \int h_0''(u)^2 du$.

Maximizing this penalized likelihood is a complex functional optimization problem. It is possible to approximate the solution on a spline basis (Joly et al., 1998):

$$\tilde{h}_0(t) = \sum_{k=1}^{m} \theta_k \psi_k(t) \tag{13}$$

where $\psi_k(t)$ is the associated spline basis $(\psi_1, \cdots, \psi_m)$ of this finite dimensional space.

The likelihood expression of $L(\beta, h_0)$ then becomes

$$L(\beta, \theta) = -\sum_{k=1}^{m} \theta_k \sum_{i=1}^{n} \exp(X_i \beta) \Psi_k(\tilde{T}_i) + \sum_{i \in \mathcal{O}} (\log h_0(\tilde{T}_i) + X_i \beta) \tag{14}$$

where the cumulative hazard function $H_0(t)$ is approximated with a proper spline basis $\Psi_k(t)$, $k = 1, \cdots, m$; such that $\Psi_k(t) = \int_0^t \psi_k(u)du$. Then, the cumulative baseline hazard function is approximated by

$$\tilde{H}_0(t) = \sum_{k=1}^{m} \theta_k \Psi_k(t) \tag{15}$$

The penalty function becomes a function of the parameters vector $\bar{\theta} = (\theta_1, \cdots, \theta_m)^\top$ and can be expressed in a matrix form as (Wahba (1983))

$$J(\theta) = \bar{\theta}^\top \Omega \, \bar{\theta} \tag{16}$$

where Ω is a symmetric $m \times m$ definite positive matrix with elements

$$\omega_{kr} = \int \psi_k''(t)\psi_r''(t)dt \tag{17}$$

We are interested in the asymptotic properties of the vector $\hat{\xi} = (\hat{\beta}, \hat{\theta})^\top$. Then, we deduce the asymptotic behavior of $(\tilde{h}_0(t) - h_0(t))$ and of $(\tilde{S}(t) - S(t))$ in order to get confidence intervals for the baseline hazard function and the survival function respectively.

Remark : why consider $(\tilde{h}_0(t) - h_0(t))$ instead of $(\widehat{h}_0(t) - h_0(t))$? The error $(\tilde{h}_0(t) - \widehat{h}_0(t))$ depends directly on the number of knots used in the spline approximation. Theoretically, one can set a sufficiently large number of knots in order to minimize the error.

4.1 Variance estimation for the survival and hazard functions

The survival and hazard functions are approximated by linear combinations of cubic splines basis. The basis $\psi_k(t)$ and $\Psi_k(t)$, $k = 1, \cdots, m$ then becomes a M-spline basis and I-spline basis respectively (Rondeau et al. (2003)) :

$$\tilde{S}(t) = \exp\Big(-H(t)\Big) = \exp\Big(-\sum_{i=1}^{m} \theta_i I_i(t)\Big) \tag{18}$$

$$\tilde{h}(t) = \sum_{i=1}^{m} \theta_i M_i(t) \tag{19}$$

For the hazard function, the variance is directly obtained from the expression (19) as a linear combination of the gaussian vector $\bar{\theta} = (\theta_1, \cdots, \theta_m)^\top$:

$$\text{var}(\tilde{h}(t)) = \bar{M}(t)^\top \text{var}(\bar{\theta}) \bar{M}(t) \tag{20}$$

where $\bar{M}(t) = (M_1(t), \cdots, M_m(t))^\top$ is the M-spline basis used in the hazard function approximation.

Then we have the following asymptotic normal distribution as $n \to \infty$:

$$\sqrt{n}\big(\tilde{h}(t, \bar{\theta}) - h(t, \bar{\theta})\big) \xrightarrow{l} N\Big(0, \bar{M}(t)^\top \text{var}(\bar{\theta}) \bar{M}(t)\Big) \tag{21}$$

The survival function case leads us to use the multivariate Delta Method (Appendix A) as the survival function is expressed by

$$\tilde{S}(t) = \exp\Big(-H(t)\Big) = \exp\Big(-\sum_{i=1}^{m} \theta_i I_i(t)\Big) \tag{22}$$

If we have the following asymptotic properties for the parameters vector $\bar{\theta}$:

$\bar{\theta} \xrightarrow{p} \bar{\theta}_0$ and $\sqrt{n}(\bar{\theta} - \bar{\theta}_0) \xrightarrow{l} N(0, \text{var}(\bar{\theta}))$ as n $\to \infty$, then, the Delta Method gives

$$\sqrt{n}\big(\tilde{S}(t,\bar{\theta}) - S(t,\bar{\theta}_0)\big) \xrightarrow{l} N\Big(0, \nabla S(t,\bar{\theta}_0)^\top \text{var}(\bar{\theta}) \nabla S(t,\bar{\theta}_0)\Big) \tag{23}$$

where

$$\begin{aligned}
\nabla S(t,\bar{\theta}_0) &= \Big(\frac{\partial}{\partial\theta^{(1)}} S(t,\theta_0^{(1)}), \cdots, \frac{\partial}{\partial\theta^{(m)}} S(t,\theta_0^{(m)})\Big) \\
&= \Big(-I_1(t)\exp(-\sum_{i=1}^{m}\theta_i I_i(t)), \cdots, -I_m(t)\exp(-\sum_{i=1}^{m}\theta_i I_i(t))\Big)
\end{aligned}$$

The confidence interval of level $1-\alpha$ for these two functions are :

$$\begin{aligned}
IC_{1-\alpha}\big(S(t,\bar{\theta}_0)\big) &= \tilde{S}(t,\bar{\theta}) \pm \Phi(1-\frac{\alpha}{2})\Big(\frac{\nabla S(t,\bar{\theta}_0)^\top \text{var}(\bar{\theta}) \nabla S(t,\bar{\theta}_0)}{n}\Big)^{1/2} \\
IC_{1-\alpha}\big(h_0(t,\bar{\theta}_0)\big) &= \tilde{h}_0(t,\bar{\theta}) \pm \Phi(1-\frac{\alpha}{2})\Big(\frac{\bar{M}(t)^\top \text{var}(\bar{\theta}) \bar{M}(t)}{n}\Big)^{1/2}
\end{aligned}$$

where Φ is the quantile function of the standard normal distribution.

5 Simulations

We performed a simulation study to compare the statistical properties in terms of coverage rates of the different variance estimators. Two cases were considered: without covariate in a simple survival analysis and with covariates through a proportional hazard model.

5.1 Inference for hazard and survival functions

5.1.1 Design

Here we are interested in the coverage rates of survival or hazard functions estimators. The methodology is to generate a sample of observations following a known distribution, then compute the associated theoretical survival or hazard function and see if the expected (known) function belongs to the estimated confidence interval of the estimated function. The coverage rates have been computed as a mean coverage rates of 100 confidence intervals taken at equally-spaced times between the smallest and highest observed times of event. Three variance estimators are compared: the Bayesian estimator, the non-penalized sandwich estimator, and the penalized sandwich estimator.

We have generated a right-censored data sample from the Weibull distribution W(a,b). We denote a the shape parameter and b the scale parameter. The density function of a Weibull distribution is given by :

$$f(x) = \frac{a}{b}\big(\frac{x}{b}\big)^{a-1} \exp\big\{ -\big(\frac{x}{b}\big)^{a}\big\} \quad , x > 0 \tag{24}$$

The parameter values have been set at $a = 13$ and $b = 100$ in order to be close to real survival times. All observations are between 55 and 105. The censoring proportion was set at 20% and uniformly distributed on the sample. The hazard function was approximated using cubic splines and the number of knots was set at 7. The smoothing coefficient was chosen by the approximated cross-valdation criterion LCV_a. The coverage rates were computed as an average over 1000 replicas. In order to see how each method is affected by sample size, the estimations were realized for three different sample sizes : 100, 500 and 1000.

5.1.2 Results

Table 1 (resp. Table 2) show the mean coverage rates obtained with the three variance estimators for the survival function (resp. the hazard function). It appears that the Bayesian estimator gives good coverage rates whatever the sample size for both survival and hazard functions. The non-penalized sandwich estimator seems not to work correctly. The penalized sandwich estimator works correctly and seems to be more accurate for the hazard function than for the survival function.

Observations	Bayes	Non Penalized Sandwich	Penalized Sandwich
n=100	92%	98%	93%
n=500	93%	95%	87%
n=1000	95%	93%	89%

Table 1: Coverage rate of confidence intervals for the survival function based on three estimators of the variance: Bayesian estimator, non-penalized and penalized sandwich estimators.

5.2 Inference for regression coefficient in a proportional hazard model

5.2.1 Design

The second simulation is a study of a proportional hazard model with covariates. Survival times have been generated from a Weibull distribution as described below (Bender et al.

Observations	Bayes	Non Penalized Sandwich	Penalized Sandwich
N=100	92%	58%	98%
N=500	93%	51%	97%
N=1000	93%	50%	97%

Table 2: Coverage rate of confidence intervals for the hazard function

(2005)). The proportional hazard model is given by

$$h_i(t) = h_0(t)\exp(X_i\beta) \quad , \quad i = 1, \cdots, n. \tag{25}$$

where X is the vector of the p covariates, β the vector of the regression coefficients and $h_0(t)$ is the baseline hazard function. Let U be a random variable with the uniform distribution $U([0,1])$. The survival times T of the proportional hazard model, following a Weibull distribution with parameters λ (scale) and k (shape) can be expressed as

$$T = \frac{1}{\lambda}[-\log(U)\exp(X_i\beta_i)]^{1/k} \tag{26}$$

The simulations were made with one and two independent covariates. The regression coefficients have been set to $\beta_1 = 1$, $\beta_2 = -1$. The number of survival times was $n = 3000$. The Weibull distribution parameters were W(12,1/0.01). As for the first simulation of section section 5.1, the parameter values of the Weibull distribution have been set so that the resulting distribution is close to that of real survival times. All observations of events were between 55 and 105. The values of covariates X_1 and X_2 were generated from uniform distributions: $X_1 \sim U([0,1])$, $X_2 \sim U([0,3])$. The censoring level was set at 20%. We focused on the Bayesian estimator. We generated 1000 replicas of these experiments.

5.2.2 Results

Table 3 shows results with one covariate. Table 4 shows results for two covariates. In both cases we see that the bias of the penalized likelihood estimator is very small and that the coverage rates based on the Bayesian estimator of the variance are satisfactory.

6 Conclusion

We have recalled facts about the penalized likelihood approach in its use for estimating an unknown function. We have focussed on the case of survival data. We have compared by

β_i	$\widehat{\beta}_i$	$\widehat{\sigma}_{Bayes}(\widehat{\beta}_i)$	CI width (Bayes)	Rate (Bayes)
$\beta_1 = 1$	1.00036	0.05	0.20	95%

Table 3: Inference with the the Bayesian estimator for regression coefficients in a proportional hazard model: 1000 simulations with 2 covariates.

β_i	$\widehat{\beta}_i$	$\widehat{\sigma}_{Bayes}(\widehat{\beta}_i)$	CI width (Bayes)	Rate (Bayes)
$\beta_1 = 1$	1.003	0.06	0.24	96%
$\beta_2 = -1$	-0.969	0.08	0.34	91%

Table 4: Inference with the the Bayesian estimator for regression coefficients in a proportional hazard model: 1000 simulations with 2 covariates.

simulation two possible estimators of the variance of the parameters, one based on the M-estimator theory, the other based on a Bayesian interpretation of penalized likelihood. We have studied coverage rates of confidence intervals for the hazard and the survival function. Both approaches work correctly, although the Bayesian estimator looks more accurate. For inference for regression coefficients we have focussed on the Bayesian estimator and found very good coverage property of the confidence intervals based on it.

7 Appendix A : Multivariate Delta Method

Let B be the estimate of a real parameter β. We assume that we have the following asymptotic properties for the estimator B, as $n \to \infty$:

$$\begin{aligned} B &\xrightarrow{p} \beta \quad \text{(consitency)} \\ \sqrt{n}\ (B-\beta) &\xrightarrow{l} N(0,\Sigma) \quad \text{(asymptotic normality)} \end{aligned}$$

where Σ is a given covariance matrix.

We have the following approximation by using a first order Taylor expansion

$$g(B) \sim h(B) + \nabla h(B)^\top (B-\beta) + o(B) \tag{27}$$

Then,

$$\begin{aligned}
\text{var}(g(B)) &\sim \text{var}(h(B) + \nabla h(B)^\top (B - \beta)) \\
&= \text{var}(h(B) + \nabla h(B)^\top B - \nabla h(B)^\top \beta)) \\
&= \text{var}(\nabla h(B)^\top B) \\
&= \nabla h(B)^\top \ \text{var}(B) \ \nabla h(B) \\
&= \nabla h(B)^\top \ \frac{\Sigma}{n} \ \nabla h(B)
\end{aligned}$$

Thus we have the following asymptotic normality

$$\sqrt{n}(g(B) - g(\beta)) \xrightarrow{p} N\Big(0, \nabla h(B)^\top \ \Sigma \ \nabla h(B)\Big) \tag{28}$$

References

Bender, R., T. Augustin, and M. Blettner (2005). Generating survival times to simulate cox proportional hazards models. *Statistics in medicine 24*(11), 1713–1723.

Chen, J. and L. Zhang (2009). Asymptotic properties of nonparametric m-estimation for mixing functional data. *Journal of Statistical Planning and Inference 139*(2), 533–546.

Cheng, G. and J. Huang (2010). Bootstrap consistency for general semiparametric m-estimation. *The Annals of Statistics 38*(5), 2884–2915.

Commenges, D., P. Joly, A. Gegout-Petit, and B. Liquet (2007). Choice between semi-parametric estimators of markov and non-markov multi-state models from generally coarsened observations. *Scandinavian Journal of Statistics 34*, 33–52.

Commenges, D., C. Proust-Lima, C. Samieri, and B. Liquet (2012). A universal approximate cross-validation criterion and its asymptotic distribution. *Arxiv preprint arXiv:1206.1753*.

Cox, D. D. and F. O'Sullivan (1990). Asymptotic analysis of penalized likelihood and related estimators. *The Annals of Statistics*, 1676–1695.

Cox, D. D. and F. O'Sullivan (1996). Penalized likelihood-type estimators for generalized nonparametric regression. *Journal of multivariate analysis 56*(2), 185–206.

Hall, P., J. Marron, and B. U. Park (1992). Smoothed cross-validation. *Probability Theory and Related Fields 92*(1), 1–20.

Huber, P. J. (1964). Robust estimation of a location parameter. *The Annals of Mathematical Statistics 35*(1), 73–101.

Joly, P., D. Commenges, C. Helmer, and L. Letenneur (2002). A penalized likelihood approach for an illness–death model with interval-censored data: application to age-specific incidence of dementia. *Biostatistics 3*(3), 433–443.

Joly, P., D. Commenges, and L. Letenneur (1998). A penalized likelihood approach for arbitrarily censored and truncated data: application to age-specific incidence of dementia. *Biometrics*, 185–194.

O'Sullivan, F. (1986). A statistical perspective on ill-posed inverse problems. *Statistical Science 1*(4), 502–518.

O'Sullivan, F. (1988). Nonparametric estimation of relative risk using splines and cross-validation. *SIAM Journal on Scientific and Statistical Computing 9*(3), 531–542.

Rondeau, V., D. Commenges, and P. Joly (2003). Maximum penalized likelihood estimation in a gamma-frailty model. *Lifetime data analysis 9*(2), 139–153.

Stefanski, L. A. and D. D. Boos (2002). The calculus of m-estimation. *The American Statistician 56*(1), 29–38.

Van der Vaart, A. (2000). *Asymptotic Statistics*. Cambridge University Press.

Wahba, G. (1983). Bayesian" confidence intervals" for the cross-validated smoothing spline. *Journal of the Royal Statistical Society. Series B (Methodological)*, 133–150.

Yu, Y. and D. Ruppert (2002). Penalized spline estimation for partially linear single-index models. *Journal of the American Statistical Association 97*(460).

Zeger, S. L. and K.-Y. Liang (1986). Longitudinal data analysis for discrete and continuous outcomes. *Biometrics*, 121–130.